

The stochastic matching problem

F. Altarelli,^{1,2} A. Braunstein,^{3,1} A. Ramezanzpour,¹ and R. Zecchina^{1,3,2}

¹*Physics Department and Center for Computational Sciences,
Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy*

²*Collegio Carlo Alberto, Via Real Collegio 30, 10024 Moncalieri, Italy*

³*Human Genetics Foundation, Torino, via Nizza 52, 10126 Torino, Italy*

The matching problem plays a basic role in combinatorial optimization and in statistical mechanics. In its stochastic variants, optimization decisions have to be taken given only some probabilistic information about the instance. While the deterministic case can be solved in polynomial time, stochastic variants are worst-case intractable. We propose an efficient method to solve stochastic matching problems which combines some features of the survey propagation equations and of the cavity method. We test it on random bipartite graphs, for which we analyze the phase diagram and compare the results with exact bounds. Our approach is shown numerically to be effective on the full range of parameters, and to outperform state-of-the-art methods. Finally we discuss how the method can be generalized to other problems of optimization under uncertainty.

One important aspect of the statistical physics approaches to disordered systems is the broad range of their interdisciplinary applications. Systems with frustration, structural disorder and uncertainties are in fact ubiquitous in many fields of science and their study has greatly benefited from the algorithms which have emerged at the interface between statistical physics of disordered systems and computer science.

One of the key problems has been the so called matching problem [1], which for the case of random instances was among the first to be solved by statistical physics methods [2] and later by rigorous mathematical techniques [3]. Matching is a constituent part of many problems in different fields, ranging from physics (dimer models [4]), to computer science (vision [5]), economics (auctions [6]) and computational biology (pattern matching [7]). It can be formulated simply (given a graph, find the largest possible subset of edges without common vertices), and is of polynomial complexity [8].

The stochastic version of matching is a basic example of *optimization under uncertainty* [9, 10], which consists in finding the minimum of a cost function depending on some stochastic parameters, given just some partial information about their value. Most real-world optimization problems involve uncertainty: the precise value of some of the parameters is often unknown, either because they are measured with insufficient accuracy, or because they are stochastic in nature and determined only *after* some decisions have been taken. The objective of the optimization process is thus to find solutions which are optimal in some probabilistic sense, a fact which introduces fundamental conceptual and computational challenges. Stochastic matching problems are in fact known to belong to higher computational complexity classes ranging from NP-hard to PSPACE-complete [8] depending on how stochasticity is introduced.

Here we apply a new method for stochastic optimization problems to the two-stage matching problem. This new method, which builds on the formalism of Survey

Propagation (SP) [11–13] and of the cavity method, is partly analytic and allows to optimize the expectation of a stochastic cost function by estimating the statistics of its minima, without resorting to explicit (and costly) sampling techniques.

In the following we define the problem, describe the method we propose for solving it, and discuss its phase diagram. We find that for large connectivity the problem enters a computationally “hard” phase where standard heuristics fail. In particular, we perform a detailed comparison with Stochastic Programming using state-of-the-art solvers. While our method has a good performance in both phases, Stochastic Programming turns out to be impractically slow in the region of large connectivity, and to have a significantly worse performance than our method in the region of small connectivity. Finally, we report about applications to problems which are NP-hard also in their deterministic setting.

The two-stage stochastic matching problem. We study a variant of the problem introduced in [14–16], where it is shown to be NP-complete. We are given a bipartite graph $G = (L, R, E)$ with L further partitioned in L_1 and L_2 , and a set of independent probabilities $\mathbf{p} = \{p_{l_2} \in]0, 1[, l_2 \in L_2\}$. The nodes in L_1 are deterministic, while the nodes in L_2 are stochastic: $l_2 \in L_2$ will be available for matching with probability p_{l_2} . In the first stage the nodes in L_1 are matched, knowing only the probabilities \mathbf{p} . In the second stage, the available nodes in L_2 are extracted according to \mathbf{p} and they are matched. The objective is to maximise the size of the final matching.

We introduce two sets of binary variables, $\mathbf{x}_1 = \{x_{l_1 r} \in \{0, 1\}, (l_1 r) \in E : l_1 \in L_1\}$ and $\mathbf{x}_2 = \{x_{l_2 r} \in \{0, 1\}, (l_2 r) \in E : l_2 \in L_2\}$, to represent the possible $M \subset E$, with $x_{l_1 r} = 1$ iff $(l_1 r) \in M$. We also introduce a set of binary parameters $\mathbf{t} = \{t_{l_2} \in \{0, 1\}, l_2 \in L_2\}$ with $t_{l_2} = 1$ iff l_2 is available for matching in the second stage. We define an energy function $\mathcal{E}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{t})$ counting the number of unmatched vertices among the available ones.

The problem consists in finding

$$\mathbf{x}_1^* = \arg \min_{\mathbf{x}_1} \mathbb{E}_{\mathbf{t}} \min_{\mathbf{x}_2} \mathcal{E}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{t}) \quad (1)$$

subject to the matching constraints $\sum_{l \in \partial r} x_{lr} \leq 1$ ($\forall r \in R$), $\sum_{r \in \partial l_1} x_{l_1 r} \leq 1$ ($\forall l_1 \in L_1$) and $\sum_{r \in \partial l_2} x_{l_2 r} \leq 1$ ($\forall l_2 \in L_2$), where $\partial r = \{l \in L : (lr) \in E\}$ and similarly for ∂l_1 and ∂l_2 .

The main difficulty of the problem stems from the fact that $\mathbb{E}_{\mathbf{t}} \min_{\mathbf{x}_2} \mathcal{E}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{t})$ has a highly non trivial dependence on \mathbf{x}_1 . In order to overcome this difficulty, we shall use the cavity method to first compute the minimum energy relative to \mathbf{x}_2 for *fixed* \mathbf{x}_1 and \mathbf{t} , and then to compute the average over \mathbf{t} of this quantity.

Minimizing relative to \mathbf{x}_2 for fixed \mathbf{x}_1 and \mathbf{t} . Once \mathbf{x}_1 is determined and the stochastic parameters \mathbf{t} are set, it is straightforward to find the optimal \mathbf{x}_2 . A possible way of doing this is by Max-Sum (MS), as discussed in [17]. We introduce the cavity fields $u_{l_2 r}$ propagating from $l_2 \in L_2$ to $r \in R$ and $h_{l_2 r}$ propagating in the opposite direction. The MS equations are $u_{l_2 r} = -\max[-1, \max_{r' \in \partial l_2 \setminus r} h_{l_2 r'}]$ and $h_{l_2 r} = -\max[-1, \max_{l'_2 \in \partial r \setminus l_2} u_{l'_2 r}]$. These equations can be solved by iteration, and their solution allows to compute $\mathcal{E}^*(\mathbf{x}_1, \mathbf{t}) = \min_{\mathbf{x}_2} \mathcal{E}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{t})$, which is found to be

$$\begin{aligned} \mathcal{E}_1(\mathbf{x}_1) - \sum_{l_2 \in L_2} \max[-1, \max_{r \in \partial l_2} h_{l_2 r}] + \\ - \sum_{r \in R} \max[-1, \max_{l_2 \in \partial r} u_{l_2 r}] + \sum_{(l_2 r) \in E: l_2 \in L_2} \max[0, h_{l_2 r} + u_{l_2 r}] \end{aligned} \quad (2)$$

where $\mathcal{E}_1(\mathbf{x}_1)$ is the energy contribution of L_1 nodes and is constant relative to \mathbf{x}_2 .

A difficulty can arise if the solution is not unique and different solutions have different energies. In this case, only one of the solutions will correspond to the actual minimum. The following argument (confirmed by numerical investigations) suggests that this is not a problem.

The MS equations are closed for cavity fields with support in $\{-1, 1\}$, and also for cavity fields with support in $\{-1, 0, 1\}$. Solutions with other supports can exist for finite size instances and for appropriate initial conditions, but we have verified numerically that they disappear in the infinite size limit, so we shall ignore them.

Let us consider (as in [17] for the non-bipartite case) the uniform ensemble of instances with poissonian degree distribution and average degree c , in the infinite size limit. The average fraction p_+^u of cavity fields $u_{l_2 r}$ that take the value $+1$ satisfies the equation $p_+^u = \exp[-c \exp(-c p_+^u)]$. Also $(1 - p_-^u)$, p_+^h and $(1 - p_-^h)$ must satisfy the same equation. In the case of bipartite graphs p_+^h can be different from p_+^u (and p_-^h from p_-^u), and this is a notable difference relative to the non-bipartite case. In any case, p_+^h and p_-^h are determined from p_+^u and p_-^u .

For $c < e$ the equation admits an unique solution, which implies that $p_+^u = (1 - p_-^u)$ and $p_+^h = (1 - p_-^h)$,

meaning that the cavity fields have support over $\{-1, 1\}$. This unique distribution of cavity fields will correspond to an *essentially* unique fixed point of MS: it is possible that some disconnected components (with finite size) admit several fixed points, but the fixed point of the $O(N)$ component (which dominates the energy) is unique. This is confirmed by numerical simulations.

For $c > e$, the situation is more complicated: the equation $x = \exp[-c \exp(-cx)]$ admits 3 solutions, and p_+^u can be different from $1 - p_-^u$ (and p_+^h from $1 - p_-^h$). The condition $p_+^u + p_-^u \leq 1$ implies that $p_+^u \leq 1 - p_-^u$, so the total number of solutions will be at most 6. Only some of these possible solutions will correspond to positive values of the energy (and the other ones can be dismissed), and only one of them will be the correct one. We have studied in detail the case for $c = 5$: the number of solutions corresponding to positive energy is 3, and remarkably the value of the energy is the same for all of them. One of the solutions has support on $\{-1, 0, 1\}$, corresponding to the 1-RSB case in the non bipartite case [17], and the remaining two have support in $\{-1, 1\}$.

On finite size instances, we have verified numerically that these 3 fixed points can always be obtained by choosing appropriate initial conditions. Their energies are close to each other, but not exactly the same, and the correct one is always the largest.

We conclude from this discussion that the energy computed from (2) is correct for instances extracted with poissonian degree distributions with $c < e$ and approximately correct for instances with $c > e$. It must be noted, however, that the reduced instance to be solved in the second stage is not necessarily poissonian, as the probability that a node in R is matched to a node in L_1 can be correlated to its degree. Moreover, it is possible that some small disconnected components have multiple solutions, that combined with the 3 solutions of the giant component give a larger number of fixed points, but these will always have energies that are approximately equal. We shall neglect these possible issues, comforted by our numerical results.

In the following we shall give the explicit computations for the case where the support of \mathbf{u} and \mathbf{h} is $\{-1, 1\}$, but not for the case where the support is $\{-1, 0, 1\}$: even though we have implemented both cases, we have verified that the energies of the solutions obtained are almost exactly the same for all connectivities; however, the expressions for case $\{-1, 0, 1\}$ are much more complicated, and the running times are much longer.

Computing the average relative to \mathbf{t} . To proceed with the computation of the average in (1), we note that the energy (2) is a sum of local terms over \mathbf{x}_1 , \mathbf{h} and \mathbf{u} , so that its average can be computed with a procedure similar to Survey Propagation (SP) [11–13]. This corresponds to a *simplified* Belief Propagation (BP) for the variables \mathbf{x}_1 , \mathbf{h} , \mathbf{u} and \mathbf{t} , where \mathbf{x}_1 , \mathbf{h} and \mathbf{u} are subject to hard constraints implementing the matching conditions

and the MS update equations, and where \mathbf{t} are subject to external fields that force them to take the marginal $p(\mathbf{t}) = \prod_{l_2 \in L_2} \mathbb{P}[t_{l_2} = 1] = \prod_{l_2 \in L_2} p_{l_2}$.

We introduce the probabilities $U_{lr} = \mathbb{P}_{\mathbf{t}}[u_{lr} = 1]$ propagating from left to right, and $H_{lr} = \mathbb{P}_{\mathbf{t}}[h_{lr} = 1]$ propagating from right to left. The SP-like equations for U_{lr} and H_{lr} are:

$$U_{lr} = p_l \prod_{r' \neq r} (1 - H_{lr'}) , \quad H_{lr} = \prod_{l' \neq l} (1 - U_{l'r}) \quad (3)$$

Equations (3) can be derived by observing that $U_{lr} = \mathbb{P}[t_l = 1] \mathbb{P}[-\max(-1, \max_{r' \neq r} h_{lr'}) = 1 | t_l = 1]$, and similarly for H_{lr} .

The average minimum energy is then computed by averaging (2) over \mathbf{t} using $U_{lr} = \mathbb{P}_{\mathbf{t}}[u_{lr} = 1]$, and $H_{lr} = \mathbb{P}_{\mathbf{t}}[h_{lr} = 1]$:

$$\begin{aligned} \mathcal{E}^*(\mathbf{x}_1) = \mathbb{E}_{\mathbf{t}} \mathcal{E}^*(\mathbf{x}_1, \mathbf{t}) = & \sum_l p_l \left[2 \prod_r (1 - H_{lr}) - 1 \right] + \\ & + \sum_r \left[2 \prod_l (1 - U_{lr}) - 1 \right] + 2 \sum_{(lr)} H_{lr} U_{lr} \quad (4) \end{aligned}$$

where the term $\mathcal{E}_1(\mathbf{x}_1)$ is included and represented with the convention $U_{lr} \equiv H_{lr} \equiv x_{lr} \forall l \in L_1$. For example, the contribution from a vertex $l \in L_2$ is +1 if the vertex is present and if all the incoming values of h_{lr} are -1, which happens with probability $p_l \prod_r (1 - H_{lr})$; the same contribution will be -1 if the vertex is present and if there is at least one incoming value of h_{lr} equal to +1, which happens with probability $p_l [1 - \prod_r (1 - H_{lr})]$. The average of the contribution is then $p_l [2 \prod_r (1 - H_{lr}) - 1]$. The average of all remaining terms is computed similarly.

Notice that the “naive” application of BP to the problem defined over the variables \mathbf{x}_1 , \mathbf{h} , \mathbf{u} and \mathbf{t} (subject to the appropriate external fields), in which one would consider the pairs (h_{lr}, u_{lr}) as single joint variables, with cavity probabilities $\mathbb{P}_{l \rightarrow r}[(h_{lr}, u_{lr})]$ and $\mathbb{P}_{r \rightarrow l}[(h_{lr}, u_{lr})]$, would lead to the *wrong* result for $c > e$. In fact, when the number of fixed points of the MS equations depends on \mathbf{t} , the naive procedure would give to each of them a weight proportional to $p(\mathbf{t})$ while the correct weight is $p(\mathbf{t})/n_{\mathbf{t}}$, where $n_{\mathbf{t}}$ is the number of fixed points corresponding to a given \mathbf{t} and $p(\mathbf{t})$ is its probability. This is achieved with the SP procedure we introduced.

Minimizing relative to \mathbf{x}_1 . We can then proceed to minimize this energy, using again MS. We consider the messages U_{lr} and H_{lr} as variables of a new problem and introduce the cavity messages $\mathcal{U}_{lr}(U) = \log \mathbb{P}[U_{lr} = U]$ propagating from left to right and $\mathcal{H}_{lr}(H) = \log \mathbb{P}[H_{lr} = H]$ from right to left. Notice that if $l \in L_1$, we will have $U_{lr} = x_{lr} \in \{0, 1\}$ satisfying the matching constraints $\sum_r U_{lr} \leq 1$, while if $l \in L_2$ we will have $U_{lr} \in [0, 1]$ satisfying the SP update equations (3), and similarly for H_{lr} . The continuous distributions over messages associated to \mathbf{x}_2 variables can be discretized for numerical purposes.




FIG. 1. Average energy density vs. average connectivity of L nodes. The three lines correspond (from top to bottom) to a greedy algorithm (assign \mathbf{x}_1 as if $\mathbf{t} = 0$), to a “smart” greedy algorithm (find the maximum-weight matching on the full instance with weights on the nodes equal to their probability to be available and assign \mathbf{x}_1 accordingly), to the SP-derived algorithm (with \mathbf{h} and \mathbf{u} with support on $\{-1, 1\}$), and to the offline lower bound of the optimum (with prior knowledge of \mathbf{t}). The vertical line is at $c = e$. Each point is an average of 50 to 100 instances, with error bars smaller than the point sizes. The instances have $|L_1| = 1000$ and $|L_2| = |R| = 2000$, with p_l distributed uniformly in $[0, 1]$.

The update equations for the messages \mathcal{U}_{lr} and \mathcal{H}_{lr} are obtained as usual for MS, i.e. $\mathcal{U}_{lr}(U_{lr}) = \max \left[-E_{lr}(U_{lr}, H_{lr'}) + \sum_{r' \neq r} \mathcal{H}_{lr'}(H_{lr'}) \right]$, where $E_{lr}(U_{lr}, H_{lr'})$ is the sum of the terms in (4) containing U_{lr} , and where the maximisation is over the values of the incoming messages $\{H_{lr'} : r' \neq r\}$ subject to the appropriate constraints; the update of $\mathcal{H}_{lr}(H_{lr})$ is obtained similarly. We don’t report these equations for brevity. All these maximizations can be performed efficiently by exploiting their associativity. In order to improve the convergence of the algorithm we also introduce a reinforcement term for the messages associated to edges $(l_1 r)$ with $l_1 \in L_1$ [18, 19].

These equations can be solved by iteration starting with uniform initial conditions. These are the *only* message passing equations that need to be solved numerically. At the fixed point, the values of \mathcal{U}_{lr} and \mathcal{H}_{lr} provide the optimal values of \mathbf{x}_1 by setting $x_{lr} = 1$ if and only if $[\mathcal{U}_{lr}(1) - \mathcal{U}_{lr}(0)] + [\mathcal{H}_{lr}(1) - \mathcal{H}_{lr}(0)] + 2 > 0$. Once \mathbf{x}_1 has been assigned and the realization of \mathbf{t} has been extracted it is easy to perform the minimization over \mathbf{x}_2 .

Numerical results and comparison with other methods. Figure 1 shows some results obtained with the SP-derived algorithm in the case where \mathbf{h} and \mathbf{u} have support on $\{-1, 1\}$. The case where \mathbf{h} and \mathbf{u} have support on $\{-1, 0, 1\}$ gives results that are very close to these.

In order to give quantitative evidence of the potential-

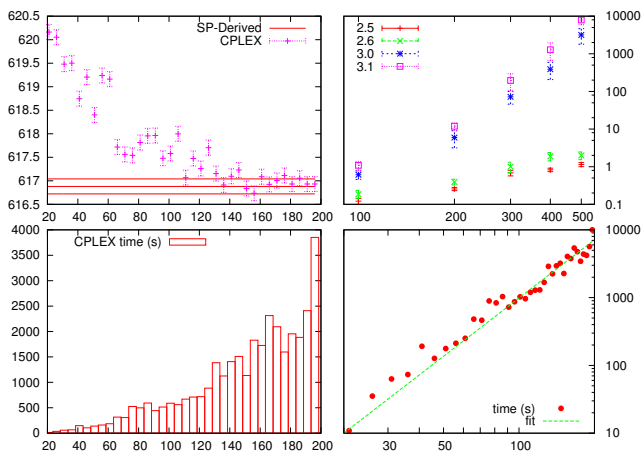


FIG. 2. Each point represents an instance with $|L_2| = |R| = 2|L_1|$ and p_i distributed uniformly in $[0, 1]$. Top Left: Energy vs. number of samples ρ obtained by CPLEX for $c = 2.5$ and $|L_1| = 1000$, compared with the one computed by the SP-derived algorithm. The energies and error bars were computed by resampling over 10 000 samples. Bottom Left: CPLEX time (seconds) vs. number of samples in the same instances. Top Right: CPLEX time as a function of $|L_1|$ for several values of c and $\rho = 10$. Bottom Right: Best fit of CPLEX times in Bottom Left with $f(x) = bx^a$ gives $a \approx 2.35$.

ities of our approach for real world problems, we have made two final studies: on the one hand we have compared the performance with state-of-the art method, and on the other we have applied the method to problems which are NP-hard even in the deterministic setting. In both cases, the results perfectly corroborate our expectations. (i) We compared the SP-derived algorithm with two other standard approaches. The first is a greedy strategy solving a weighted matching based on $p(\mathbf{t})$: even though it is very fast, its solutions are much worse (Figure 1). The second is called stochastic programming. It consists in extracting ρ realizations $\mathbf{t}^1, \dots, \mathbf{t}^\rho \sim p(\mathbf{t})$ and then solving $\min_{\mathbf{x}_1} \sum_{i=1}^{\rho} \min_{\mathbf{x}_2^i} \mathcal{E}(\mathbf{x}_1, \mathbf{x}_2^i, \mathbf{t}^i) = \min_{\mathbf{x}_1, \mathbf{x}_2^1, \dots, \mathbf{x}_2^\rho} \sum_{i=1}^{\rho} \mathcal{E}(\mathbf{x}_1, \mathbf{x}_2^i, \mathbf{t}^i)$ using OR techniques like linear relaxations complemented with branch-and-bound. Note that this minimization problem is NP-Complete[14]. We employed two well known tools for this task: iLog CPLEX, a commercial, industrial strength linear/integer programming software from IBM, and `lp.solve`, an open source alternative. Although qualitatively similar, results with `lp.solve` were uniformly worse than the ones of CPLEX, so we will not report them. We observe that the results depend strongly on ρ and on the average degree c . As expected, for fixed c the quality of the solution improves as ρ increases, but the running time becomes larger. For c up to around 2.5, CPLEX seems to be able to solve the problem in polynomial time in both ρ and N , but either it is much slower than the SP-derived algorithm or it gives a significantly higher energy (de-

pending on ρ). For $c = 3.5$ and above, the time scaling of CPLEX worsens significantly: for $\rho = 10$, the running time increases dramatically with $|L_1|$, and for $|L_1| = 1000$ CPLEX was not able to attain an optimum under a cut-off of 24 hours even for $\rho = 2$. Note that the SP-derived algorithm employs around one minute. (ii) The method was also successfully applied to a stochastic version of the maximum weight independent set problem [20], when the node's contribution to the total weight is uncertain but its distribution is known. This is a relevant problem in communication networks with some interference constraints [21]. Details will be given in [22].

RZ acknowledges the EU grant n. 265496.

-
- [1] L. Lovasz, M.D. Plummer, *Matching Theory*, North-Holland (1986)
 - [2] M. Mezard and G. Parisi, *J. de Physique* **48**, 1451 (1987)
 - [3] D. Aldous, *Random. Struct. Algor.* **18**, 381 (2001) .
 - [4] R. J. Baxter, *Exactly solved models in statistical mechanics* Academic Press, London (1982).
 - [5] E. Trucco, A. Verri, *Introductory techniques for 3D computer vision* Prentice All (1998)
 - [6] D. P. Bertsekas, *A Distributed Algorithm for the Assignment Problem* Lab. for Information and Decision Systems Report, MIT (1979)
 - [7] D. Gusfield, *Algorithms on strings, trees, and sequences: computer science and computational biology* Cambridge University Press (1997)
 - [8] C. Papadimitriou, *Computational complexity* J. Wiley (2003)
 - [9] A. Prekopa, *Stochastic programming* Springer (1995)
 - [10] J. Birge and F. Louveaux, *Introduction to stochastic programming* Springer Verlag (1997)
 - [11] M. Mezard and R. Zecchina, *Phys. Rev. E* **66**, 56126 (2002)
 - [12] A. Braunstein and R. Zecchina, *J. Stat. Mech.* (2004), P06007
 - [13] A. Braunstein, M. Mezard, and R. Zecchina, *Random Structures and Algorithms* **27**, 201 (2005)
 - [14] N. Kong and A. Schaefer, *Eur. J. Oper. Res.* **172**, 740 (2006)
 - [15] I. Katriel, C. Kenyon-Mathieu, and E. Upfal, *Theoretical Computer Science* **408**, 213 (2008)
 - [16] B. Escoffier, L. Gourves, J. Monnot, and O. Spanjaard, *Eur. J. Oper. Res.* **205**, 19 (2009)
 - [17] L. Zdeborova and M. Mezard, *J. Stat. Mech.* (2006), P05003
 - [18] A. Braunstein and R. Zecchina, *Phys. Rev. Lett.* **96**, 30201 (2006)
 - [19] F. Altarelli, A. Braunstein, J. Realpe Gomez, and R. Zecchina, *J. Stat. Mech.* (2009), P07002
 - [20] S. Sanghavi, D. Shah, and A. S. Willsky, *IEEE Transactions on Information Theory*, **55** (11), 4822 (2009).
 - [21] D. Shah, *Performance Modeling and Engineering*, Springer US, 147-184 (2008).
 - [22] F. Altarelli, A. Braunstein, A. Ramezanzpour, and R. Zecchina, *Statistical physics of optimization under uncertainty*, in preparation (2011)